COLT 2021 Virtual-only Days

August 4, 2021

7:00–7:10	**Program Chairs' Welcome Remark**
7:10–12:00	**Learning Theory Alliance Mentoring Workshop Session 1**
12:00–18:00	No program
18:00–18:30	Networking
18:30 – 21:30	**Tutorial A** **Statistical Inference in Distributed and Constrained Settings** *Jayadev Acharya, Clément Cannone, Himanshu Tyagi*
21:30 – 22:30	Social Hour

*All times are in MT (Mountain Time)

August 5, 2021

7:00–7:15	**Networking**
7:15 – 10:15	**Tutorial B** **Statistical Foundations of Reinforcement Learning** *Akshay Krishnamurthy, Wen Sun*
10:15–10:45	Break and Social Games
10:45 – 11:45	**Remembering Matthew Brennan** *Guy Bresler*
11:45–12:00	Break
12:00 – 13:30	**WiML-T Event** *Career Panel and Speed Networking*
13:30–18:00	No program
18:00–23:00	**Learning Theory Alliance Mentoring Workshop Session 2**

*All times are in MT (Mountain Time)

LeT-All Mentoring Workshop

The workshop will have three components: academic workshop, technical talks and social events. The academic workshop will have two tracks: (1) undergrad/junior-grad track focused on grad school admissions and fellowships, and (2) senior-grad/postdoc track focused on academic jobs. The technical talks will feature 3 spotlight talks (15 minutes each) covering recent exciting work in learning theory. The social component will include speed-networking, mentoring tables, games and more.

Session 1 speakers include Boaz Barak, Dylan Foster, Satyen Kale, Emilie Kaufmann, Katrina Ligett, Ankur Moitra, Shay Moran, Ariel Procaccia, Cynthia Rush and Rocco Servedio.

Session 2 speakers include Suriya Gunasekar, Kevin Jamieson, Anna Karlin, Tengyu Ma, Praneeth Netrapalli and Yisong Yue.

See detailed schedule here.

Tutorial A

Statistical Inference in Distributed or Constrained Settings: Techniques and Recipes

Jayadev Acharya, Clément Cannone and Himanshu Tyagi

The goal of this tutorial is to provide the attendees with an overview of techniques and recipes for distributed estimation and testing under constraints. Over the recent years, many papers have obtained both upper and lower bounds for statistical estimation under communication, local privacy, and memory constraints: these questions are motivated by applications in machine learning and distributed computing, and are at the intersection of theoretical computer science, machine learning, statistics, and information theory. Our goal is to provide a primer of those techniques, aiming to provide both an understanding of the underlying challenges and ideas, and "plug-and-play" general recipes the attendees could then apply to the problems of their choice. Our focus will be on establishing lower bounds for statistical estimation, in particular for parameter estimation (single- and high-dimensional) and testing. The tutorial will cover various models: nonadaptive, sequentially adaptive, blackboard model, and memory-constrained settings; with applications to high-dimensional parameter estimation and testing. (If time allows, we will briefly cover extensions to nonparametric density estimation.)

[**60 min** Tutorial: Part 1] Setting the stage: scope, models, and lower bounds for learning and estimation problems.

[**45 min** Recitation-style practice/clarification session] Applications to Gaussian mean estimation under local privacy or communication constraints; derivation of optimal bounds, comparison between existing methods.

[**45 min** Tutorial: Part 2] Lower bounds for testing problems.

Tutorial B

Statistical Foundations of Reinforcement Learning

Akshay Krishnamurthy and Wen Sun

The past decade has seen tremendous interest in sequential decision making under uncertainty, a broad class of problems involving an agent interacting with an unknown environment to accomplish some goal. Reinforcement learning approaches to addressing these problems have led to recent AI breakthroughs in game playing, robotics, and elsewhere. Inspired by these empirical demonstrations, many researchers from the learning theory community have turned their attention to reinforcement learning, in an attempt to better understand these problems and develop new algorithmic principles. Their efforts have led to a more-modern statistical foundation for reinforcement learning that places an emphasis on non-asymptotic characterizations via global convergence, sample complexity, and regret analyses.

This tutorial will provide an overview of this emerging theory with a focus on the most challenging online exploration setting. The tutorial is organized into three sections:

[**45 min** Tutorial: Part 1] will cover the necessary background and definitions. We focus here on the most basic setting of tabular Markov Decision Processes and consider problems of increasing difficulty: from planning, to optimization with an exploratory distribution, to online exploration. We will present two algorithms: Natural Policy Gradient for the optimization problem and UCB-Value Iteration for exploration, along with their guarantees.

[**45 min** Recitation-style practice/clarification session] will be organized in smaller discussion groups led by the tutorial presenters and other researchers in the field. We will cover the analyses for NPG and UCB-VI in detail, highlighting key lemmas that are broadly useful in reinforcement learning, as well as technical connections to related fields.

[**60 min** Tutorial: Part 2] will focus on online exploration beyond the tabular setting, where function approximation is required for generalization. Here we will provide a tour of the zoo of RL models and complexity measures that enable tractable learning, as well as some statistical barriers. We will close with some open problems and future directions.

The tutorial will be accessible to all COLT attendees. No background knowledge in RL is required, but we do expect tutorial attendees to be comfortable with the standard mathematical tools used in learning theory research, such as concentration inequalities and some linear algebra.